UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/749,405	12/28/2000	Stephan J. Jourdan	2207/7085	5261

25693 7590 11/21/2006

KENYON & KENYON LLP
RIVERPARK TOWERS, SUITE 600
333 W. SAN CARLOS ST.
SAN JOSE, CA 95110

EXAMINER	
LI, AIMEE J	

ART UNIT	PAPER NUMBER
2183	

DATE MAILED: 11/21/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 09/749,405
Filing Date: December 28, 2000
Appellant(s): JOURDAN ET AL.

Sumit Bhattacharya (Reg. No. 51,469)
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 30 December 2005 appealing from the Office action mailed 31 May 2005.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

Tran, U.S. Patent Number 5,822,575

Heuring, Vincent P. and Jordan, Harry F. Computer Systems Design and Architecture. Reading, Mass.: Addison Wesley Longman, Inc., ©1997. Page 537.

Rosenberg, Jerry M. Dictionary of Computers, Information Processing, and Telecommunications. Second Edition. New York: John Wiley & Sons, ©1987. Page 632, "tag".

McFarling, Scott. "WRL Technical Note TN-36: Combining Branch Predictors". Palo Alto, CA: Digital Western Research Laboratory, ©1993. Pages 1-25.

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claims 24-26 are rejected under 35 U.S.C. 102(b) as being taught by Tran, U.S. Patent Number 5,822,575 (herein referred to as Tran).

Referring to claim 24, Tran has taught a method for restoring a branch prediction apparatus following a branch misprediction of a branch instruction, comprising:

Restoring a base misprediction history register (Tran column 14, line 14 to column 15, line 7; column 18, lines 44-62; column 19, lines 31-49; Figure 3; and Figure 4); and

Restoring a branch predictor history register (Tran column 14, line 14 to column 15, line 7; column 18, lines 44-62column 19, lines 31-49; Figure 3; and Figure 4).

Referring to claim 25, Tran has taught updating a branch predictor (Tran column 14, line 14 to column 15, line 7; column 18, lines 44-62column 19, lines 31-49; Figure 3; and Figure 4).

Referring to claim 26, Tran has taught updating a meta predictor (Tran column 14, line 14 to column 15, line 7; column 18, lines 44-62column 19, lines 31-49; Figure 3; and Figure 4).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 1-8, 10-15, and 17-23 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chang, Hao, and Patt's "Alternative Implementations of Hybrid Branch Predictors" (herein referred to as Patt) in view of McFarling's "WRL Technical Note TN-36: Combining Branch Predictors" (herein referred to as McFarling).

Referring to claim 1, Patt has taught a branch prediction apparatus, comprising:

A meta predictor to receive as inputs an index value and a branch prediction to generate a misprediction value in accordance with said inputs (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2); and

A logic gate to receive said branch prediction and said misprediction value to generate a final prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Patt has not taught a base misprediction history register to provide an output. McFarling has taught a base misprediction history register providing an output (McFarling page 12, paragraph 2). In regards to McFarling, the counter is similar to the misprediction history register, since it contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling page 12, lines 13-14 and Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the misprediction history register of McFarling in the device of Patt.

Referring to claim 2, Patt has not taught wherein said base misprediction history register includes misprediction history data. McFarling has taught wherein said base misprediction history register includes misprediction history data (McFarling page 12, paragraph 2). In regards to McFarling, the counter is similar to the misprediction history register, since it contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling page 12, lines 13-14 and Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the misprediction history register of McFarling in the device of Patt.

Referring to claim 3, Patt has taught an instruction that provides said index value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 4, Patt has taught wherein said instruction is a branch instruction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2)

Referring to claim 5, Patt has taught wherein said final prediction determines a branch for said branch instruction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 6, Patt has taught a branch predictor that receives said index value and generates said branch prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 7, Patt has taught wherein said branch predictor utilizes a prediction scheme to generate said branch prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 8, Patt has taught wherein said branch predictor includes a target address field and a prediction table (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 10, Patt has taught a method for predicting branches, comprising:

> Receiving an index value and a branch prediction value correlating to said index value at a meta predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2); and

Generating a misprediction value at said meta predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Patt has not taught a misprediction history value. McFarling has taught a misprediction history value (McFarling page 12, paragraph 2). In regards to McFarling, the counter is similar to the misprediction history register, since it contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling page 12, lines 13-14 and Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the misprediction history register of McFarling in the device of Patt.

Referring to claim 11, Patt has taught generating said branch prediction value at a branch predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 12, Patt has taught receiving an index value at said branch predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 13, Patt has taught generating a final prediction according to said branch prediction and said misprediction value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 14, Patt has taught determining a final value, and updating said meta predictor and said base misprediction history register according to said final value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 15, Patt has taught wherein said updating includes comparing said final value to said branch prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 17, Patt has taught a processor, comprising:

> A branch predictor to generate a branch prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2);
>
> A meta predictor to receive an index value, said branch prediction data to generate a misprediction value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Patt has not taught a base misprediction history register. McFarling has taught a base misprediction history register (McFarling page 12, paragraph 2). In regards to McFarling, the counter is similar to the misprediction history register, since it contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling page 12, lines 13-14 and Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the

invention was made to incorporate the misprediction history register of McFarling in the device of Patt.

Referring to claim 18, Patt has taught a final prediction to correlate to said misprediction value and said branch prediction value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 19, Patt has taught a logic gate to generate said final prediction (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 20, Patt has taught a computer readable medium having stored a plurality of executable instructions, the plurality of instructions comprising instructions to:

> Receiving an index value and a branch prediction value correlating to said index value at a meta predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2); and
>
> Generating a misprediction value at said meta predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Patt has not taught a misprediction history value. McFarling has taught a misprediction history value (McFarling page 12, paragraph 2). In regards to McFarling, the counter is similar to the misprediction history register, since it contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling

page 12, lines 13-14 and Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the misprediction history register of McFarling in the device of Patt.

Referring to claim 21, Patt has taught an instruction to generate said branch prediction value at a branch predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 22, Patt has taught an instruction to receive an index value at said branch predictor (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Referring to claim 23, Patt has taught an instruction to generate a final prediction according to said branch prediction and said misprediction value (Patt page 252, column 2, paragraph 4; page 255, column 1, paragraph 2 to column 2, paragraph 3; and Figure 2).

Examiner would like to note, in response to the "Order Returning Undocketed Appeal to Examiner", that claim 27 was originally rejected in the Non-Final Office Action (dated 18 March 2004). The rejection of claims 24-27 was maintained in the second Non-Final Office Action (dated 19 November 2004) and Final Office Action (dated 31 May 2005). The Examiner made a typographical error in the second Non-Final Office Action, Final Office Action, and Examiner's Answer by accidentally omitting the rejection of claim 27. However, Applicants did recognize the rejection of claim 27. The Applicant has grouped claim 27 with independent claim 24, from which claim 27 depends. The grouping of claim 27 with claim 24 is evidenced in Applicant's own responses. Applicant's remarks on 18 August 2004 argued the allowability of claims 24-27

based upon limitations found in claim 24 starting on page 7. Applicant's remarks on 21 March 2005 makes no mention of claim 27, however, the remarks maintain the arguments in regard to claim 24. Applicant's Appeal Brief filed on 30 December 2005 does not argue claim 27 individually and states on page 13 of the Appeal Brief "Claims 25-27 depend from the aforementioned allowable independent claims, and therefore are in condition for allowance as well." That is the only mention of claim 27 within the entirety of the Appeal Brief and Reply Brief filed 05 June 2006. As can be seen by these responses, Applicant recognized that claim 27 rises and falls with claim 24, since the limitation being argued is a limitation found in claim 24. The following was the rejection of claim 27 from the Non-Final Office Action (dated 18 March 2004) and maintained through the history of the application.

Claim 27 is rejected under 35 U.S.C. 103(a) as being unpatentable over Tran, U.S. Patent Number 5,822,575 (herein referred to as Tran), as applied to claim 24 above, in view of Applicant's admitted prior art (herein referred to as Prior Art). Tran has not taught flushing an instruction pipeline processing said branch instruction. Prior Art has taught flushing an instruction pipeline processing said branch instruction (Prior Art page 2, lines 16-17). A person of ordinary skill in the art at the time the invention was made would have recognized that flushing the pipeline after a mispredicted branch eliminates incorrect instructions and data, thereby ensuring the system is not contaminated with incorrect instructions and data. Therefore, it would have been obvious to a person of ordinary skill in the art at the time this invention was made to incorporate flushing the pipeline after a mispredicted branch of Prior Art in the device of Tran.

The following is a table that more clearly maps the most comprehensive independent claims with the prior art cited and explains how the prior art falls within the meets and bounds of the broadest reasonable interpretation of the claim language.

Instant Application	Prior Art Rejection and Explanation
Claim 1	
A branch prediction apparatus, comprising:	
A meta predictor to receive as inputs an index value and a branch prediction to generate a misprediction value in accordance with said inputs; and	Patt Page 252, column 2, paragraph 4 Page 255, column 1, paragraph 2 to column 2, paragraph 3 Figure 2 Applicant defines "meta predictor" on page 5, lines 4-5 as "meta predictor 104 *is* a misprediction predictor". Applicant does not have an explicit and deliberate definition of "misprediction value" anywhere in the specification. The closest definition is on page 6, lines 25-26 which states "Misprediction value 112 then *may be used* to decide whether to reverse the prediction provided by the base predictor, or branch prediction 108." However, because of the "may be used" language, this definition was not considered explicit and deliberate. Therefore, the Examiner interpreted

	"misprediction value" as any type of value that reflects branch mispredictions in some manner. Whether the "misprediction value" reversed the prediction or not was considered an improper reading of limitations from the specification into the claim language. Patt has taught on page 255, column 1, paragraph 2 to column 2, paragraph 3 that the Branch Predictor Selection Table (BPST) tracks which branch prediction, e.g. branch prediction from the claim, is the most accurate for a given branch instruction, e.g. the index value in the claim (Patt page 255, column 2, lines 6-11). The BPST then uses the branch instruction address and the current branch history to output a value reflecting which branch prediction is more accurate, e.g. misprediction value in the claim since it reflects which value has less of a chance at mispredicting the branch, that a multiplexer uses to determine which prediction to use (Patt page 255, column 2, lines 12-15). A simplified version of this device is shown in Figure 2, where the BPST outputs a value to the prediction multiplexer to choose between prediction 1,

	P1, and prediction 2, P2, to output a final prediction, Prediction.
A logic gate to receive said branch prediction and said misprediction value to generate a final prediction.	Patt Page 252, column 2, paragraph 4 Page 255, column 1, paragraph 2 to column 2, paragraph 3 Figure 2 The logic gate is the multiplexer receiving the two predictions and BPST output to determine the final prediction. As stated above, the misprediction value is the value from the BPST, since it reflects which prediction is more accurate, i.e. which one has the least likelihood of mispredicting the branch instruction. The branch prediction is either P1 or P2, whichever the BPST reflects is more accurate. The final prediction is the prediction, P1 or P2, output by the multiplexer as the Prediction.
A base misprediction history register to provide an output	McFarling Page 12, paragraph 2 Patt has not explicitly taught "a branch misprediction history register to provide an output" that affects the misprediction value. However, as stated above, Patt has taught that the BPST "records which predictor was most frequently correct for the times this branch occurs

	with the associated branch history (Patt page 255, column 2, lines 6-9)" and that it is a counter (Patt page 255, column 2, lines 14-15), but Patt has not disclosed the specifics on the BPST and how it tracks the accuracy. McFarling has taught that there are multiple counters, which are a form of registers (For more information please see Heuring and Jordan's Computer Systems Design and Architecture ©1997 page 537, Section A.16 Counters), reflecting the past history of the branches, since the counters are saturated up or down depending upon whether a branch is predicted or mispredicted. The counter is a misprediction history register, since it is a register that contains data related to past mispredictions. A person of ordinary skill in the art at the time the invention was made would have recognized that the counter tracks which predictor is more accurate, thereby ensuring the combined branch predictor picks the best predictor to use and reducing mispredictions which cost time (McFarling page 12, lines 13-14 and

	Patt page 255, column 2, lines 6-9 and 15-18). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the misprediction history register of McFarling in the device of Patt.
Claim 24	
A method for restoring a branch prediction apparatus following a branch misprediction of a branch instruction, comprising:	
Restoring a base misprediction history register; and	Tran Column 13, lines 42-65 Column 14, line 14 to column 15, line 7 Column 18, lines 44-62 Column 19, lines 31-49 Figure 3 Figure 4 In the specification, the closest to a explicit, deliberate definition for "base misprediction history register" is on page 5, lines 19-20 which states "...base misprediction history register 110 reflects the correctness of the base predictor standing alone..." and page 5, lines 21-22 which states "...base misprediction

	history register 110 records whether previous branch predictions were correctly predicted by the base predictor." However, due to the context of these statements and the discussion of the exemplary embodiment in the specification, it was unclear whether this was meant as a definition of the term "base misprediction history register" or just an analysis of its function within the environment of the exemplary invention. Tran has taught two elements that could be considered the base misprediction history register. The first, Tran describes in column 13, lines 17-35 a register storing branch tags indicating valid branches being predicted. Tran describes in column 14, lines 36-38 has a misprediction tag which is used to show whether the branch has been mispredicted or not. In column 14, lines 42-54, Tran describes the branch shift register is updated accordingly to reflect whether the branch was mispredicted or not. Specifically, the branch shift register purges data relating to the incorrectly predicted instructions, so that the data

	reflects about whether past branches were mispredicted or not. The second, Tran describes in column 13, lines 50-59 a valid indication register which indicates whether a branch in the pipeline is valid or not, e.g. whether the branch prediction information is correct or incorrect. This reflects whether the branch prediction information was correctly predicted or mispredicted, so it is a register that contains the misprediction history of the current branches in the pipeline. In column 15, lines 48-54 resets the valid bits in the valid register 60 when a branch has been mispredicted. Tran also teaches in column 14, lines 55-65 that the storage 56, which includes the valid register 60, needs to be adjusted to reflect a branch misprediction.
Restoring a branch predictor history register.	Tran Column 13, lines 42-65 Column 14, line 14 to column 15, line 7 Column 18, lines 44-62 Column 19, lines 31-49 Figure 3 Figure 4 Tran describes in column 13, lines 43-45 a taken/not-taken register that stores whether the branch

	instructions are predicted taken or not taken. Tran also teaches in column 14, lines 55-65 that the storage 56, which includes the valid register 60, needs to be adjusted to reflect a branch misprediction. A further description of this type of register is in column 18, lines 2-4. Tran describes in column 18, lines 44-46 that the taken/not-taken register needs to be recovered to the pattern prior to the mispredicted branch.

The following is a table showing how the limitations in claim 10, 17, and 20 relate to the limitations of claim 1, so the rejections and explanations of the rejections for the limitations in claim 1 in the above table apply accordingly. The limitations similar to each other are either bolded, underlined, or italicized to reflect their similarities.

Claim 1	Claim 10	Claim 17	Claim 20
A branch prediction apparatus, comprising:	A method for predicting branches, comprising:	A processor, comprising:	A computer readable medium having stored a plurality of executable instructions, the plurality of instructions

			comprising instructions to:
A **base misprediction history register** to provide an output;		**A base misprediction history register**;	
A meta predictor to receive as inputs *an index value and a branch prediction* to generate a misprediction value in accordance with said inputs and **said base misprediction history register output**; and	Receiving *an index value, a branch prediction value correlating to said index value*, and **a misprediction history value** at a meta predictor; and Generating a misprediction value at said meta predictor.	A *branch predictor to generate a branch prediction*; A meta predictor to receive *an index value, said branch prediction* and **base misprediction history register data** to generate a misprediction value.	Receiving *an index value, a branch prediction value correlating to said index value*, and **a misprediction history value** at a meta predictor; and Generating a misprediction value at said meta predictor.
A logic gate to receive said branch prediction and said misprediction value			

to generate a final prediction.			

(10) Response to Argument

Applicants' argue in essence on pages 7-8

> ... Therefore, this Patt counter and any value associated with does not actually lead to or factor into any prediction, but merely keeps track of an overall accuracy of a group of predictors. In other words, any value associated with the counter does not have any active, determinative purpose (like for example, a misprediction value), but rather is merely used to passively reflect the overall effectiveness of the prediction scheme. Therefore, Applicants submit that this "counter" and any associated counter cannot be the equivalent of a "meta predictor to receive as inputs an index value and a branch prediction to generate a misprediction value"...

This has not been found persuasive. Patt describes on page 255, column 2, lines 12-18 that the associated counter outputted by the BPST is used to more accurately select which branch prediction to use. Patt also shows this in Figure 2. The BPST outputs to the multiplexer to select whether to use branch prediction P1 or branch prediction P2 as the final prediction, e.g. the output labeled Prediction. This is how the counters factor into the overall branch prediction scheme. As stated in the explanation table of the rejection above, Patt has taught on page 255, column 1, paragraph 2 to column 2, paragraph 3 that the Branch Predictor Selection Table (BPST) tracks which branch prediction, e.g. branch prediction from the claim, is the most accurate for a given branch instruction, e.g. the index value in the claim (Patt page 255, column 2, lines 6-11). The BPST then uses the branch instruction address and the current branch history to output a value reflecting which branch prediction is more accurate, e.g. misprediction value in the claim since it reflects which value has less of a chance at mispredicting the branch, that a

multiplexer uses to determine which prediction to use (Patt page 255, column 2, lines 12-15). A simplified version of this device is shown in Figure 2, where the BPST outputs a value to the prediction multiplexer to choose between prediction 1, P1, and prediction 2, P2, to output a final prediction, Prediction.

Applicants' argue in essence on pages 8-9 and 10

> The Patt reference does not contain any such reference, teaching or suggestion to the ability to reverse the prediction provided by a base predictor.

This has not been found persuasive. As stated in the rejection explanation table above, Applicants do not have an explicit and deliberate definition of "misprediction value" anywhere in the specification. The closest definition, which is argued in this Appeal brief, is on page 6, lines 25-26 which states "Misprediction value 112 then *may be used* to decide whether to reverse the prediction provided by the base predictor, or branch prediction 108 (emphasis added)." However, because of the "may be used" language, this definition was not considered explicit and deliberate. The "may be used" language not only suggests that this is an example from the language "may be", since "may be" reflects possibility not definitiveness. Also, the language reflects a use for the misprediction value, not a definition of what the value is. Therefore, the Examiner interpreted "misprediction value" as any type of value that reflects branch mispredictions in some manner. Whether the "misprediction value" reversed the prediction or not was considered an improper reading of limitations from the specification into the claim language.

Applicants' argue in essence on page 9

> Applicants respectfully disagree and submit that these limitations were not meant to be read in to the claims...Any such explanation taken from the specification is intended to describe and clarify one functionality of the 'misprediction value', and to further clarify that any such functionality is not disclosed in the cited sections.

The Examiner finds this statement contrary to the arguments presented in the previous arguments and in the Appeal Brief. For example, the Appeal Brief blatantly states on page 9, lines 13-14 "The Patt reference does not contain any such reference, teaching or suggestion to the ability to reverse the prediction provided by a base predictor", which effectively summarizes the arguments up to this point with regards to the term "misprediction value". Applicant wishes the functionality of the term "misprediction value" described within the specification to be considered a distinguishing point between the instant application and the prior art used to reject the claims, and uses this distinction to show how the prior art does not read upon the current claim language. However, it is unclear how such a functionality can distinguish the claimed invention from the prior art when no such functionality is claimed, but located within the specification, without being read into the claims from the specification, as reflected in Applicants' own arguments in the Appeal Brief when stating on page 9, lines 7-12

> Indeed, Applicants submit there is no disclosure of any value that remotely resembles a misprediction value as described in embodiments of the present application. By way of background, further support for this limitation can be found at least at page 6 line 23 of specification which states:
>
> > Misprediction value 112 then may be used *to decide whether to reverse the prediction provided by the base predictor*, or branch prediction 108.

None of this language is recited within the claims. Specifically, the "to decide whether to reverse the prediction provided by the base predictor" is not located within the claim language but in the specification, and Applicants' arguments responded to above are with regard to how this distinguishes the instant application from the prior art.

Applicants argues in essence on pages 10-11 and 12-13

> The Office Action asserts that the branch tag shift is the equivalent of the 'base misprediction history register', as the register's data is changed based upon the misprediction tag. Applicants disagree. First, as is well known in the art a tag is an instruction that specifies how a document is to be formatted (e.g., as in HTML). In light of this, Applicants submit that the formatting aspect inherent in the operation of this branch tag shift of Tran has no resemblance to a base misprediction history register.
> ...Clearly, this supports Applicant's assertion that the branch tags of Tran are separate from an actual prediction. Furthermore, and more importantly, Applicants submit such a branch tag shift does not disclose a base misprediction history *register*...

This has not been found persuasive. The "tag" being referred to in the arguments is that found in webpage design. Applicant is correct that a "tag" in the webpage design environment formats how a document looks, such as tags in HTML format text, spacing, etc. within the webpage. However, "tag" in the computer architecture environment has another meaning. As it is defined in the computer architecture area, a "tag" is a set of characters that represent information associated with a set of data (For more information see Jerry Rosenberg's Dictionary of Computers, Information Processing, and Telecommunications Second Edition ©1987). In this instance, and as supported by its use in Tran, "tag" represents information with regards to the branch instructions. There is no "formatting aspect inherent in the operation of this branch tag shift" as Applicants' arguments state. As stated in the explanation of the rejection above, Tran has taught two elements which could be considered the base misprediction history register. The first, Tran describes in column 13, lines 17-35 a register storing branch tags indicating valid branches being predicted. Tran describes in column 14, lines 36-38 has a misprediction tag which is used to show whether the branch has been mispredicted or not. In column 14, lines 42-54, Tran describes the branch shift register is updated accordingly to reflect whether the branch

was mispredicted or not. Specifically, the branch shift register purges data relating to the incorrectly predicted instructions, so that the data reflects about whether past branches were mispredicted or not. The second, Tran describes in column 13, lines 50-59 a valid indication register which indicates whether a branch in the pipeline is valid or not, e.g. whether the branch prediction information is correct or incorrect. This reflects whether the branch prediction information was correctly predicted or mispredicted, so it is a register that contains the misprediction history of the current branches in the pipeline. In column 15, lines 48-54 resets the valid bits in the valid register 60 when a branch has been mispredicted. Tran also teaches in column 14, lines 55-65 that the storage 56, which includes the valid register 60, needs to be adjusted to reflect a branch misprediction. Tran states in column 13, line 19 that element 50 is a branch tag *register* and in column 13, lines 43 that element 58 is a taken/not-taken *register*. With regards to the global shift register and as shown in the rejection explanation table above, this was equated to the branch predictor history register, since it records whether the branches were predicted taken or not-taken.

Applicant argues in essence on pages 11-12

> Therefore, in order to be a proper rejection, Applicants submit the Tran reference must disclose a *base* misprediction history register...

As stated in the explanation table of the rejection above, the closest to a explicit, deliberate definition for "base misprediction history register" is on page 5, lines 19-20 which states "...base misprediction history register 110 reflects the correctness of the base predictor standing alone..." and page 5, lines 21-22 which states "...base misprediction history register 110 records whether previous branch predictions were correctly predicted by the base predictor."

However, due to the context of these statements and the discussion of the exemplary embodiment in the specification, it was unclear whether this was meant as a definition of the term "base misprediction history register" or just an analysis of its function within the environment of the exemplary invention. The arguments emphasize "base", however, there is nothing in the specification showing significance to this distinction other than that cited in the previous sentences. Tran has taught two elements which could be considered the base misprediction history register. The first, Tran describes in column 13, lines 17-35 a register storing branch tags indicating valid branches being predicted. Tran describes in column 14, lines 36-38 has a misprediction tag which is used to show whether the branch has been mispredicted or not. In column 14, lines 42-54, Tran describes the branch shift register is updated accordingly to reflect whether the branch was mispredicted or not. Specifically, the branch shift register purges data relating to the incorrectly predicted instructions, so that the data reflects about whether past branches were mispredicted or not. The second, Tran describes in column 13, lines 50-59 a valid indication register which indicates whether a branch in the pipeline is valid or not, e.g. whether the branch prediction information is correct or incorrect. This reflects whether the branch prediction information was correctly predicted or mispredicted, so it is a register that contains the misprediction history of the current branches in the pipeline. In column 15, lines 48-54 resets the valid bits in the valid register 60 when a branch has been mispredicted. Tran also teaches in column 14, lines 55-65 that the storage 56, which includes the valid register 60, needs to be adjusted to reflect a branch misprediction. In both cases, the branch tag register and the valid indication register reflect whether the branch was mispredicted or not in its stored data.

Applicants argue in essence on page 12

> Applicants respectfully disagree and submit that these limitations were not meant to be read in to the claims. As argued above, the claim language, specifically the term "base misprediction history register", more than sufficiently describes an operation...Any such explanation taken from the specification is intended to describe and clarify one functionality of the "misprediction value" (*sic*), and to further clarify that any such functionality is not disclosed in the cited sections.

First, the Examiner believes that Applicants intended to mean "base misprediction history register" rather than "misprediction value" in the last sentence quoted. Second, the Examiner finds this statement contrary to the arguments presented in the Appeal Brief. For example, the Appeal Brief blatantly states on page 12, lines 1-2 "Therefore, in order to be a proper rejection, Applicants submit the Tran reference must disclose a *base* misprediction history register as found in embodiments of the present application", which effectively summarizes the arguments up to this point with regards to the term "base misprediction history register". Applicant wishes the functionality of the term "base misprediction history register" described within the specification to be considered a distinguishing point between the instant application and the prior art used to reject the claims, and uses this distinction to show how the prior art does not read upon the current claim language. However, it is unclear how such a functionality can distinguish the claimed invention from the prior art when no such functionality is claimed, but located within the specification, without being read into the claims from the specification, as reflected in Applicants' own arguments in the Appeal Brief when stating on page 11, line 20 to page 12, line 2

> By way of background, support for the "base misprediction history register" can be found, among other places, at page 6 lines 16-20 of the specification, which state:
>
> > As discussed above, base misprediction history register 110 reflects the correctness of the base predictor standing alone. Unlike global history registers that record whether previous branches were taken or not taken,

> base misprediction history register 110 records whether previous branch predictions were correctly predicted by the base predictor.
>
> Therefore, in order to be a proper rejection, Applicants submit the Tran reference must disclose a *base* misprediction history register as found in embodiments of the present application.

None of this language is recited within the claims. Specifically, the "records whether previous branch predictions were correctly predicted by the base predictor" is not located within the claim language but in the specification, and Applicants' arguments responded to above are with regard to how this distinguishes the instant application from the prior art.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

Aimee J. Li

16 November 2006

Conferees:

Eddie Chan

EDDIE CHAN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100

Donald Sparks

DONALD SPARKS
SUPERVISORY PATENT EXAMINER